# EFFECTS OF MANDATED DRIVER PAY INCREASES ON THE GIG ECONOMY: EVIDENCE FROM NEW YORK CITY

Dmitri Koustas
University of Chicago

James Parrott
The New School

Michael Reich
UC Berkeley

This Draft: June 10, 2020

*Abstract*: The advent of app-dispatched services, such as Uber and Lyft, has dramatically changed urban transportation markets in recent years. The rapid increase of ridership on these services has raised concerns about the pay and working conditions of the drivers, as well as increased congestion and pollution costs in metropolitan areas. In this paper, we analyze a major regulatory change that New York City implemented in 2019 to increase driver pay by mandating minimum per minute and per mile pay rates. We make use of a large, administrative dataset at the driver-trip level containing half a billion app rides originating in New York City from August 2017 to December 2019. Our research design uses the fact that the pay mandate binds differentially across times of the day and days of the week, and across geographic areas, depending on pre-existing pay and pricing policies by the app companies and supply and demand patterns in the period before the policies were implemented. Our findings provide insight on key structural parameters specific to the industry, including labor supply and demand elasticities, and should help better inform regulation of app-based dispatch services as well as the broader "gig economy."

Acknowledgments: We thank the Sloan Foundation for generous financial support. We thank Yalun Su, Nasha Virata and Xingxing Yang for excellent research assistance. This research uses anonymized data from the New York City Taxi and Limousine Commission (TLC). We are grateful for institutional support from Bill Heinzen, Fausto Lopez, Rodney Stiles and Celine Zakrayan at the TLC. This analysis reflects the independent findings of the authors. Two of the authors (Parrott and Reich) have worked as paid consultants for the TLC in the past but were not under contract at the time of this research. We have no further relevant financial interests to disclose.

## 1. Introduction

New jobs in the app-based gig economy have exploded in recent years, with the largest concentration in app-hail For-Hire-Vehicle (FHV) transportation services. In the New York City metro area, for instance, the number of online platform workers, which includes app-based FHV services, went from 0.01 to 1.3 percent of the workforce between 2012 and 2016 (Collins et al. 2019). In 2019, app-based FHV trips averaged around 700,000 per day in New York City alone, compared with approximately 250,000 traditional taxi trips. While the growth of the industry represents a triumph of technology, expanding service and convenience to millions of riders, the rapid growth has also raised concerns about congestion and labor market conditions. Drivers are mostly independent contractors and are not subject to employment protections covering the traditional workforce. Cities are also grappling with consequences for the incumbent taxi industry, which is subject to more stringent regulation that puts them at a competitive disadvantage.

These concerns have led municipalities to experiment with local regulation of the app-dispatched FHV industry. In this paper, we study the effects of a major regulatory change, which New York City implemented in February 2019, to mandate minimum pay per mile and per minute rates. These mandated pay increases were intended to increase the hourly wage to the equivalent of New York City's $15 minimum wage, after accounting for independent contractor taxes and driver expenses (Parrott and Reich 2018). The extent to which these fare increases were effective in increasing driver pay in practice is an empirical question, depending on pass-through of the pay increase into prices, the demand response and driver labor supply responses.

The standard approach for policy evaluation is to conduct a differences-in-differences analysis, comparing outcomes pre and post a reform, relative to a control group unaffected by the reform. In the present setting, this approach is complicated because the entire city was affected by the policies and policies were bundled (in addition to a newly implemented congestion charge, a cap on new drivers implemented the previous July had begun to bind by the time the new policies went into effect). In theory, possible control groups might include adding another city that had no change in regulation; however, data are not readily available for other cities, not to mention that New York City's size and unique characteristics make finding a comparable control city difficult. Nor is the traditional ride-hail taxi industry an ideal control: taxis were also affected by the congestion charge, and would be indirectly affected by any changes in the app-hail industry to the extent that app-hail and ride-hail taxi services are substitutable.

Our approach exploits how the policy binds differently across routes, defined by times of the day, days of the week and geographic areas. The key to our research design is the dynamic nature of pricing in the app industry, which results in substantial variation in prices across routes, depending on supply and demand patterns. For instance, prior to the implementation of the new policies, 80 percent of trips within Downtown Brooklyn on Saturday at midnight were paid above the level mandated by the new pay standards; the average expected increase in pay for such trips was only about 2 percent. In contrast, 80 percent of trips from Manhattan to Central Brooklyn on Wednesday at 8pm were paid below the new pay standard; the average expected increase in pay per trip on this route was around 19 percent. In addition, only trips passing through core Manhattan were subject to a new congestion charge, allowing us to separate out its effect.

Our main empirical strategy thus compares outcomes on the same route (time of day, day of week, start and end location), before and after the policy. We exploit the *expected* increase in pay based on the formulaic increase from the pay policy.

We find that the pay standard increased average driver pay per trip by about 15 percent, compared to the 10 percent increase that the new pay standard itself predicts. The pay increase appears to be mostly passed on to riders in the form of higher fares. For a 1 percent increase in fares due to the pay standard, we estimate the implied demand elasticity to be around -0.68 percent. Our estimates of demand elasticities are largely consistent with the literature.

Our results are persistent for the 11 months for which we have data following the implementation of the policy. We also examine two sources of heterogeneity, by income and by access to public transit. We find that demand elasticities are much larger for lower-income census tracts and for those with more access to public transit.

The app-based driver workforce is heterogeneous, including a large proportion who drive full-time and a large proportion who are only casually attached to the industry (Parrott and Reich 2018). In future work, we intend to examine data on individual drivers before and after implementation of these policies by calculating each driver's exposure to the pay standard and congestion charges, based on their typical driving patterns.

New York City provides a unique laboratory to study this industry and to assess how well the new policies have fared. The recent policy actions constitute the most far-reaching regulatory measures applied to the app-dispatched car service sector in the U.S. Very importantly for our research, no other city collects such comprehensive data on the industry as a whole, linking drivers across multiple app-dispatching platforms and even to the traditional taxi sector. Most other studies of the industry use data from one particular platform and so are able only to follow driver activity on the particular app (e.g. Angrist et al. 2018, Chevalier et al. 2019, Cohen et al. 2018, Hall et al. 2019).

We proceed as follows: We briefly discuss related literature in the next section. We present the details of the policies we study in Section 3. We next describe the data and provide some descriptive findings in Section 4, before laying out our main empirical strategy in Section 5. We present results in Section 6. Section 7 concludes.

## 2. Related Literature

The most directly relevant literature pertains to the labor supply of app-dispatched drivers. A major challenge to studying the industry is the availability of data, as the new growth in the industry does not appear to be picked up by traditional government surveys of the labor market (Abraham et al 2019, Koustas 2019). As a result, the literature has turned to study the industry using non-traditional datasets.[1]

Uber's administrative data has been used in a number of studies. In one notable study, Hall and Krueger (2018) utilized administrative data on Uber drivers nationally and for selected large cities and proprietary survey data from surveys conducted in December 2014 and November 2015 to describe the driver workforce as one primarily seeking flexible work opportunities to provide supplemental income with most drivers already owning a car. Parrott and Reich (2018) use an earlier vintage of TLC data used in the present paper to study app-drivers in NYC. While the TLC are only for trips that start in NYC, a key advantage of the data is that drivers can be linked across all the app-dispatch companies, providing a more complete picture of driving behavior.

[1] Other related papers include Cohen et al. (2018); Katz and Krueger (2019); Abraham, Haltiwanger, Sandusky and Spletzer (2018); Bernhardt and Thomason (2017); Collins, Garin, Jackson, Koustas, and Payne (2019); and Farrell and Greig (2019).

The closest study to ours is the working paper by Hall, Horton and Knoepfle (2018). In this paper, the authors examine the effects of changes in the base pay for Uber drivers. A key finding is that after a change in the fare, the market primarily clears on utilization. For instance, if pay increases, this encourages more drivers to be on the road. At the same time demand falls, resulting in lower utilization (rides per driver per hour). One major limitation of the Hall, Horton and Knoepfle (2018) study is that these changes in fares only apply to Uber and that the authors only see activity on the Uber platform. Unlike in the TLC data, the authors do not observe driver and passenger substitution to other platforms, namely Lyft, which may introduce biases in measuring driver utilization. This is also a different setting than the one we propose to study, where in our case increases in driver pay apply to *all* app-based companies.

The labor market literature on taxi drivers and the gig economy is also relevant to this project. Using related administrative TLC data on the city's traditional ride-hail taxi drivers, Farber (2015) has studied their labor supply in detail. However, his work predates the rise of app-based dispatch services.

Another strand of the gig economy literature has focused on the merits of gig flexibility over more traditional forms of employment (Angrist, Caldwell and Hall 2018; Chevalier et al 2019; Koustas 2018). We do not address the debate over these costs and benefits in this paper, but instead focus on policy evaluation of the recent reforms. We do, however, intend to study effects on part-time versus full-time drivers and new versus long-run incumbent drivers.

## 3. Background and Policy Details

This section briefly describes the app-dispatched FHV market and the policies affecting the app-dispatched FHV industry over our period of study.

*Background on the app-based FHV industry*

In our period of study, the New York City market was served by four main app-based dispatch services. In order of market share, these companies are: Uber, Lyft, Juno and Via. Juno ceased operations in November 2019, leaving three companies. Drivers for Uber, Lyft and Juno are independent contractors, while Via drivers typically are hourly employees (Via's market share is very low). The companies set prices and remunerate drivers based on proprietary algorithms.

*Passenger Fares*

Passenger fares are dynamically priced. Each company uses its own proprietary algorithm that takes into account expected and actual supply of drivers and passenger demand. Prices can vary at any moment. In addition, the companies may offer passenger incentives, such as fare discounts. Tips can be voluntarily provided by passengers to drivers, on top of the base fare.

*Driver Pay*

Prior to the February 2019 pay standard, each company set its own *minimum* base fare and per minute and per mile rates. After the implementation of the pay standard in February 2019, minimum per minute and per mile rates followed the pay standard. With the exception of Via, which pays drivers an hourly rate for most rides, driver pay is dynamic, varying based on a company's proprietary algorithms. Passenger fares can be decoupled from driver pay; drivers do not always get a fixed percentage of the fare a passenger pays–before passenger incentives. In addition, different driver incentives are provided by the companies, such as driver bonuses. Via has a hybrid model. Most of its drivers are paid an hourly rate. Outside of set shifts, additional earnings opportunities are available that follow a model similar to the other FHV companies.

*Pay Standard*

The TLC pay standard affected rides that originate in the five boroughs; it went into effect on February 2, 2019. The pay standard was first announced in December of 2018, with an implementation date of January 1, but was later delayed to February 2. The pay standard mandated per mile and per minute minimums that varied by FHV platform. Uber, Lyft and Juno were mandated the following rates: $1.088 per mile in the five boroughs, $1.262 for miles accrued out of town, $0.495 per minute within the 5 boroughs, and $0.574 per minute outside of the five boroughs. The following rates were mandated for Via: $0.914 per mile in the five boroughs, and $0.416 per minute within the 5 boroughs, with the same out of time per minute and per mile rates as the other companies. Rides that start outside of the city and end in the city are not subject to the pay standard. On February 1, 2020, these rates were increased by 1.4 percent (in line with the Consumer Price Index).

The pay standard also includes a provision to reduce driver pay per trip if a company improves its utilization rate—the percentage of time drivers have a passenger in their vehicles. This adjustment would occur periodically (such as every quarter or year) and would be based on the data the companies provide to the Taxi and Limousine Commission, the regulatory agency for the industry. This provision is intended to improve the use of the drivers' time and vehicles, which were vacant nearly half of every hour. Drivers would share in the efficiency gains, as rides per hour would increase, thereby increasing driver pay per hour. The companies can control the utilization rate by managing the number of new drivers in their systems. However, these incentives to maintain utilization have yet to be implemented in practice.

*Congestion Surcharge*

A congestion surcharge went into effect on February 1, 2019, one day before the pay standard. The congestion surcharge is applied per trip, and is $2.75 for non-shared trips and $0.75 for shared rides. A congestion surcharge does not apply to wheelchair-access vehicle (WAV) trips. The congestion surcharge applies to trips that start, end or pass through core Manhattan south of 96th Street (the "Congestion Zone").

*Moratorium on New FHV Drivers*

New York City also enacted a moratorium (cap) on new FHV vehicle licenses, effective in August 2018. This policy capped non-WAV FHV vehicles at around 85,000, curtailing entry of new vehicles and drivers into the market since August 2018. However, the companies were able to increase the number of licensed vehicles by about a six-month amount between the passage of the moratorium in July 2018 and its implementation in August 2018.

*Limit on intake of new drivers by the app companies*

The companies continued to add new drivers to their systems through April 2019. At that time, Uber and Lyft each announced that they were no longer admitting new drivers into their systems. In the fall of 2019, these companies also began to manage the number of drivers who could log on to their apps at low-demand periods.

## 4. Data

The data we use are drawn from TLC-mandated data sharing agreements with FHV operators. Data have been made available to us from August 2017 through December 2019. These data cover the universe of trips with pickups in the five boroughs of New York City. A trip that starts out of town and drops off in the city *is not* included in the data. In-town pickups with out-of-town dropoffs are included in the data; approximately 8 percent of trips have an out-of-town dropoff. One notable omission from the data are trips that start at Newark airport and end in the city. Trip-level data include pickup and dropoff dates and times. Pickup and dropoff longitude and latitude are mostly available over the whole period. Shared trips in some cases are indicated directly, but otherwise can be inferred if the pickup from a next customer occurs before the drop off of the previous customer.

Some key data elements are only available in certain months. Driver pay, including tips broken out in a separate field, and passenger fare data, are only available from August 2017-June 2018, and again from February 2019, leaving a gap in pay data between July 2018-January 2018. Our understanding is that the TLC simply did not request these data from the companies for this period; the earlier data were collected for calibrating the pay standard, and the later data were collected for the purposes of enforcement. Out-of-town minutes and miles were only reported to the TLC from February 1, 2019, since they are important inputs into the pay standard. The congestion surcharge amount per trip is available in the data beginning February 2019. Additional data include timestamps for app log ins and log outs, and separate files with driver bonuses/incentives.

We next show some basic descriptive statistics from the data. Figure 1 plots the total number of trips, number of drivers, average driver pay per trip and average base fares per trip over the period for which data are available. As shown in Figure 1(a), trips are increasing dramatically in 2017, but the market appears to be relatively mature by the middle of 2018 as growth in the number of trips per week flattens out. Trips decline sharply in December 2019; part of this is seasonal, but the seasonal effect seems stronger than in 2018. Trips appear to subsequently recover however, with no sharp change in February, the month the pay standard and congestion charges are implemented.

Figure 1(b), reports total unique drivers across all platforms active (with at least one trip) over the course of the week. The figure shows that the number of drivers is increasing close to linearly in the first half of the data, by about 2,500 drivers per month, before slowing in January 2019, and subsequently declining from the second half of 2019. One interpretation of the slowing growth and subsequent decline is that the moratorium on new vehicle licenses begins to bind.

Figure 1(c) shows average driver pay and base passenger fares *per trip* when these data are available. Note that these are raw statistics, unadjusted for route characteristics such as duration or miles. We see that driver pay per trip was trending upward in February to June 2018, which could be a seasonal effect since we see a similar pattern in 2019. Overall, average driver pay per trip appears to have increased in 2019 compared with 2018. At the same time fares appear to remain low from February-May 2019, although fares rise again with time. Two major app companies were preparing to become publicly listed companies around this time, so the fare pricing decisions might be driven by competitive concerns by the app companies about losing ridership.

## 5. Empirical Strategy

Figure 1 provided an aggregate look at the FHV market. At an aggregate level, it is difficult to discern, let alone separate, the overall effects of any policy changes from other factors affecting the market. In the next section, we outline an empirical strategy making use of the rich microdata on trips. Our approach exploits variations in the treatment intensity of the pay standard across platform, times of the day, days of the week, and pickup/dropoff geography.

The pay standard mandates a single minimum fare per minute and per mile. In practice, dynamic pricing set by the app platforms means that the pay standard will bind differentially depending on pre-period prices and driver pay. As discussed above, each app company follows in its own pricing and pay policies. Prices are dynamic. Each company uses its own proprietary algorithm that takes into account expected supply of drivers and passenger demand. These prices can vary by the minute, but there is also a predictable component. With the exception of Via, which pays drivers an hourly rate for most rides, driver pay is also dynamic.

To give a sense of this variation, Figure 2(a) plots average driver pay per trip for every hour in the week, averaged over the pre-period August 2017-June 2018. Pay typically peaks in the morning rush hour, declines in the late morning and early afternoon, before rising somewhat again later in the day, but not again reaching the morning peak. Some of these differences in average pay per trip fare are due to differences in the miles and minutes of a trip, that vary due to passenger demand patterns. Others are due to differences in each company's algorithm and commission rates.

To better understand these patterns, Figures 2(b) shows the number of completed trips, and hours the drivers have the app on, by hour of day and day of week. At first glance, these supply and demand patterns mirror each other closely. First, note morning rush hours with the highest fares are also a time with comparatively fewer drivers working. Completed trips drop by more around midday than the hours drivers have the app on. This is a time when we expect fares to be relatively low. The opposite pattern can be seen on Saturday around midnight, where there is a reduction in app hours but no comparably large reduction in trips, suggesting fares will move higher.

We use this variation, and also the trip characteristics across the days and times, to construct a counterfactual expected dollar increase in pay due the pay standard for the period before the implementation of the pay standard. Define the minimum pay under the pay standard for trip *i*:

$$pay\ standard_i = \ 1.088\ in\ town\ miles_i + 1.262\ out\ of\ town\ miles_i + 0.495\ in\ town\ min_i + 0.574\ out\ of\ town\ min_i \tag{1}$$

where the coefficients on in and out of town miles follow the mandate of the pay standard. As discussed above, out-of-town minutes and miles are only available after the pay standard is implemented. To calculate counterfactual prices, we impute out-of-town minutes and miles based on Poisson regressions using trip characteristics available pre-policy: minutes, total miles and pickup-dropoff matched pairs.

For a trip that would have been paid less than the pay standard in the period before the implementation of the pay standard, define the counterfactual expected increase in pay in log points (approximately the percent increase), as follows:

$$\widehat{pay\ inc}_i = (\ln(pay\ standard_i) - \ln(driver\ pay_i)) \times \mathbb{1}(pay\ standard_i - driver\ pay_i > 0) \tag{2}$$

where $driver\ pay_i$ is the actual pay for trip $i$. If driver pay for a trip already exceeds the pay standard, then the pay standard does not bind, which is captured by the indicator in this equation that takes on the value of 1 if the pay standard binds and 0 otherwise. If the pay standard binds, $\widehat{pay\ inc}_i$ takes on the difference (in log points) between the pay standard mandate and the amount actually paid out.

In practice, the pay standard also applies for shared trips, where the driver is paid based on the total minutes and miles over all shared segments. In the data, total miles for a shared trip are not reported in the trip files prior to February 2018, so we calculate $\widehat{pay\ inc}_i$ only for non-shared trips.

To provide a better sense of the variation, Figure 3(a) shows the share of trips in the pre-period over August 2017-June 2018 that the pay standard would be expected to bind, for every hour in the week. The pay standard binds for the most trips from the late morning through the early afternoon, binding for around 68% of trips.

This counterfactual log point increase (approximately a percentage increase) is shown in Figure 3(b). Across platforms, the average expected increase ranges from 0.04 log points on Monday at midnight, to nearly 14 percent on Thursday afternoons. The overall average expected increase in pay per trip is around 10%.

To give a sense of the geographic variation, Figure 4 shows the average pay increase for two selected times of day, midnight and noon, which are roughly the least and most affected times of day, respectively, by TLC taxi pickup and dropoff zones. The figure shows considerable heterogeneity in the bite of the pay standard across locations, even within a particular hour of the day.

## a. Main Estimating Equation, Route-level Variation

In this section we introduce our main estimating equation. As discussed above, raw data are at the individual trip level. We calculate the expected increase in driver pay or fares for a particular "route". We define a route as a trip made at a particular day of the week and time of the day, between pickup and dropoff locations. For our baseline definition of a route, we assign a trip to an hour and day block based on the time of pickup. For geographic units, the five boroughs are a natural starting point for pickup geographies. Because airports have their own supply and demand patterns, we separate out the two airports in the city, JFK, and LGA. We also separate Manhattan above and below 96th Street (the border of the "Congestion Zone"), Northern Brooklyn from the rest of Brooklyn, and Long Island City from the rest of Queens. This yields 10 pickup geographies. We add in Newark airport and a general "out of town" zone for dropoff locations, yielding 12 dropoff geographies. This yields (10*12=) 120 matched pickup-dropoff pairs, and (24*7*120=) 20,160 total routes.

We also consider a second definition of a route with more precise pickup/drop-off areas: pickup and drop-off latitudes and longitudes rounded to the nearest two digits. This is approximately a 1 mile by 1 mile grid for pickup and dropoff locations. The advantage of this definition is that route characteristics will be more similar in a tighter grid. The disadvantage of moving to narrower geographies is that there will be many routes with 0 rides between routes.

Using only data from the period *before* the pay standard and congestion charges went into effect, we aggregate our counterfactual increase in pay to the route level, $r(i)$:

$$\hat{p}_{r(i)} = \frac{1}{N_r}\sum_{i\in r} \widehat{pay\,inc}_i \qquad (3)$$

In words, $\hat{p}_{r(i)}$ is the average expected increase (in log points) in driver pay on route $r(i)$ that is expected from the pay standard.

Our baseline estimating equation is given as follows:

$$y_i = \alpha_{r(i)} + \alpha_{w(i)} + \sum\nolimits_{j\in J} \left[\beta_j \mathbb{I}\{year \times month(i) = j\} \times \hat{p}_{r(i)}\right] + \epsilon_i \qquad (4)$$

where $y_i$ is a trip-level outcome of interest, $\alpha_{r(i)}$ is a fixed effect for the route, $w(i)$ indexes the week trip $i$ occurred, $\alpha_{w(i)}$ is a time fixed effect and $post_{w(i)} = \mathbb{I}\{w(i) > Week\ of\ February\ 1\ 2019\}$ is an indicator for the period after February 1 when the policy is in effect. $j$ uniquely indexes a year-month. $\alpha_{r(i)}$ captures any constant level differences in pay across routes. $\alpha_{w(i)}$ captures any common week-to-week movements in pay across all routes, such as seasonality. $\beta_j$ is our main coefficients of interest, reporting the response of $y_i$ to a log point increase in $\hat{p}_{r(i)}$, and allowing this effect to vary flexibly over time. For identification, one year-month must be omitted, and the estimated $\beta_j$ coefficients will then be relative to this month.

This specification has two key advantages. First, this specification provides a series of placebo estimates prior to February 2019, which is the month the pay standard took effect. Our empirical specification relies on a parallel-trends assumption that routes more or less affected by the pay standard would have been trending similarly in the absence of the policies. Significant $\beta_j$ estimates earlier than this date suggest a pretrend that would violate our identification assumptions. The most compelling results should show a structural break in our $\beta$ estimates beginning in February 2019. Second, a key research question is whether the policy has a persistent effect or fades away over time as supply and demand patterns adapt. By estimating separate coefficients for each month, we are able to capture these dynamics.

Our key outcomes of interest include (all in logged values): driver pay, driver pay plus tips, base passenger fare, the ratio of driver pay to base passenger fare (a measure of labor share),[2] total trip time, and passenger wait time. Since these outcomes are in log values, the interpretation of $\beta_j$ is as an elasticity. Given our variation is at the route level and because standard errors are likely to be correlated within a week, we cluster standard errors at the route and week levels.

[2] This is related to the effective commission, $commission = 1 - pay/fare$. Around a quarter of rides lose money, presumably because of passenger incentives, resulting in a negative measured *commission* when pay exceeds the fare. Because we take logs of our outcome measures, log $commission$ would be undefined. We instead examine $pay/fare$ since this will always be positive.

We also consider a more parsimonious specification, estimating one "post" coefficient:

$$y_i = \alpha_{r(i)} + \alpha_{w(i)} + \beta \hat{p}_{r(i)} \times post_{w(i)} + \epsilon_i \qquad (5)$$

The results for this estimation equation will be sensitive to the time frame used for the pre-period. If the data are trending differently by our treatment or affected by other policies that are correlated with our treatment, then this violates a parallel trends assumption implicit in our research design required for identification. Our dynamic specification provides a test of the parallel trends assumption, as we can see if the coefficients are trending differently before the policy went into effect.

We can also consider a version of our specification *within* driver, replacing $\alpha_{r(i)}$ with $\alpha_{r(i)\times d(i)}$, where $d(i)$ is a function denoteing the driver for trip $i$. In this specification, we are comparing earnings for the *same* driver on the *same* route in the pre and post period. This specification will control for different types of drivers taking on different routes between the pre and post period (such as drivers for premium services such as "Uber Black"); a disadvantage is that it requires drivers to be in the sample in the pre and post periods, which may introduce some attrition bias. We will get some idea of these biases by comparing estimates with these driver-route fixed effects with our baseline estimates.

We also test for heterogeneous effects by route characteristics. We consider the following sources of heterogeneity: (1) the average annual income in the census block, and (2) the average share of the population that commutes to work via public transit. Both of these measures come from the American Community Survey 5-year estimates through 2017. We calculate terciles of each measure across all pickup census blocks. Denote $tercile^k_{b(i)}$ as the tercile for measure $k$ ($k \in \{\text{income, public transit share}\}$) for a trip $i$ that picks up in census block $b$. We redefine a route so that a pickup geography is also within a tercile (i.e., when considering income, the Brooklyn borough will be further stratified into low, medium and high income areas). Our estimating equation allowing for heterogeneity is given by:

$$y_i = \alpha_{r(i)} + \alpha_{w(i)} + \sum_{j=1}^{3} \left[ \delta_j \mathbb{1}\{tercile^k_{b(i)} = j\} \times post_{w(t)} + \beta_j \mathbb{1}\{tercile^k_{b(i)} = j\} \times post_{w(i)} \times \hat{p}_{r(i)} \right] + \epsilon_i \qquad (6)$$

Note that the regressions in Equations (4)-(6) will yield the same estimates if we were to aggregate the data to the route by week level. For instance, (5) becomes:

$$y_{rw} = \alpha_r + \alpha_w + \beta p_r \times post_w + \epsilon_{rw} \qquad (7)$$

where *r* indexes route and *w* indexes week, and weighting the regression by the number of trips made on the route in the week *w*. In practice, we will estimate the regression on aggregated data in this way for computational convenience, with the exception of the specification with driver-by-route fixed effects. The aggregated version of the specification allows us to estimate an additional outcome: the total number of trips made on the route.

One limitation of our approach is that behavioral changes by passengers and drivers are likely to result in spillovers across treatment intensities. Results should be interpreted as the relative difference inclusive of behavioral changes across treatment intensity in the new post-policy equilibrium. Consider the following scenario: more drivers are drawn to work in an area where pay is higher relative to areas where pay is lower. If prices and pay adjust dynamically this should eliminate the relative pay differences across space, and bias our regression results towards finding *smaller* relative treatment effects from the pay standard for an outcome like driver pay, even if average pay across the two areas is increasing.

In practice, it will likely take some time for passengers and drivers to adjust their behavior. We estimate dynamic responses to see if behavior changes over time. Second, scope for behavioral responses may be more limited. While drivers have agency over where they start a shift, it is largely random where they are dispatched over the course of a shift.

Another limitation of our design is that labor supply responses are more difficult to measure. We have data for when an app is on, but we only know when a driver is active on a platform; we cannot see *where* a driver is active, unless a trip occurs. Some outcomes are a better fit for our second approach examining driver-level outcomes, discussed below.

Finally, our week fixed effect controls for any variation common across all hours and days within a week. The period we study saw the introduction of a moratorium on new vehicle licenses and a congestion charge. The congestion charge can be controlled for by adding an interaction between post and an indicator for passing through the Congestion Zone of core Manhattan. Regarding the moratorium, one important question is whether the cap binds in the same way across all hours of the day or day of the week. If it binds in the same way, then our approach of removing a time fixed effect is appropriate. However, if it binds differentially, the cap is a potentially important omitted variable. HourXtime FE and day of weekXtime FE should pick up some of this, and we can test robustness to adding these additional FE.

## b. Main Estimating Equation, Driver-level Variation

**Note: We do not include results yet in the paper, but we present our empirical strategy for comment.**

A key stated motivation for the driver pay standard was to increase hourly pay (and ultimately total take-home pay) for drivers. Estimates from the specifications discussed so far only provide insight at the route-level. Changes in supply and demand patterns could affect overall driver pay, even if certain routes see increases in pay. To address this, we calculate expected driver exposure to the pay standard based on typical driving behavior *prior to* the implementation of the policies. Again letting $d$ denote driver, we calculate the expected exposure to the pay standard:

$$\hat{p}_d = \sum_r w_{dr}^{pre} \hat{p}_r \qquad (8)$$

Where $w_{dr}^{pre}$ is the share of time in the period before the policies that the driver works on route $r$. Our driver-level specification is given by:

$$y_{d,t} = \alpha_d + \alpha_t + \beta \hat{p}_d \times post_t + \epsilon_{d,t} \qquad (9)$$

Our key outcomes of interest include hourly earnings, weekly earnings, weekly hours, and driver utilization (the share of hours worked spent with passengers).

As we did at the route level, we can similarly examine heterogeneity and dynamic effects. For heterogeneity, we examine differences by full-time versus part-time (average weekly hours above below 20 hours per week).

## 6. Results

### a. Route-level Variation

We report the results for our main specification given in Equation (4) graphically in Figure 5, and also report the regressions with a single "post" coefficient in Table 1, which we also overlay on the figure. All specifications include week and route fixed effects. In addition to our main coefficient of interest, we include an interaction term for passing through core Manhattan and being in the post period, to control for the congestion surcharge. Figure 5 plots the $\beta_j$ coefficients for select outcomes. We omit February 2018 when running (2), so coefficients are relative to one year before the policies went into effect. We calculate $\hat{p}_{r(i)}$ using data only from July 2017-January 2018 and aggregating over all platforms proportionately based on their share of trips. We then run the regressions using data from February 2018, so that no data used to construct $\hat{p}_{r(i)}$ is also used our estimation sample to avoid any mechanical correlations.

Panel (a) shows the results for log total trips, which is one of the few outcomes we have available over the entire period without any gaps. We find no strong evidence of a pre trend in our placebo periods prior to February 2020. We find total trips fall by about 1% for every 1% increase in predicted pay from the pay standard, and this fall is persistent.

Panel (b)-(d) plot pay per trip, pay per trip inclusive of tips, and the base fare, respectively. Recall that these outcomes are unfortunately not available for July 2018-January 2019. Pay per trip is an important outcome. In a regression of log pay per trip on the expected percent increase in pay per trip, the coefficient should be 1 if pay increased by as much as predicted from the pay standard. However, as shown in the figures and reported in Table 1, pay per trip increases by somewhat more than can be predicted by the pay standard: a 1 percent predicted increase in pay due to the pay standard actually increases pay by about 1.5 percent. We do not know precisely why we underpredicted the increase, but suspect it might be due to a disappearance of very low fares.

Examining fares, we see that essentially all of the increase in pay is passed through to fares. Given an average expected increase in pay of 10 percent, these estimates imply pay per trip and fares in NYC increased by about 15 percent due to the pay standard. Assuming the reduction in trips is the result of a fall in demand from the increase in fares, we can use Columns (1) and Column (4) of Table 1 to approximate a demand elasticity: fares increased by 1.536 percent, resulting in a -1.037 percent decline in trips, implying an estimated demand elasticity of -1.037/1.536=-0.68.

The results for trips, pay and fares are very persistent over the 11 month period following the introduction of the policy. This suggests supply and demand patterns do not quickly adapt to eliminate differences across time and space.

Panel (e) shows that wait times declined in the first 6 months, although this fades away after about 4 months. There is no notable change in passenger trip times (Fig 5 Panel f, Table 1 Column 7) or miles (Table 1, Column 8), which suggests that trip characteristics *within a route*, as we have defined it, did not substantially change.[3]

[3] We have also run our regressions using a narrower definition of a route with pickup and dropoff locations defined as a 1 mile X 1 mile grid, and find our results are largely unchanged.

While hourly earnings is not an outcome measure that makes sense in these route-level regressions, we can estimate the change in hourly earnings implied by our estimates. Consider the following relationship:

$$Hourly\ Pay = Pay\ Per\ Trip\ \times Trips\ Per\ Hour$$

which implies:

$$\Delta Ln(Hourly\ Pay) = \Delta Ln(Pay\ Per\ Trip) + \Delta Ln(Trips\ Per\ Hour)$$

Using our estimates reported in Table 1, the average increase in hourly pay from a 10 percent increase in pay per trip is $(1.482 * 0.1\ - 1.037 * 0.1 =)$ 4.5 percent.

We next discuss test for heterogeneity in our treatment effects by income of the pickup area, and by the public transit access (proxied by the share of workers who commute via public transit). Table 1b shows the results estimating a separate treatment effect for each income tercile. Tercile 1 refers to census blocks with average annual household income (in 2017 inflation-adjusted dollars) up to $64,000, and Tercile 3 starts at household income above $99,000. Tercile 1 is the omitted tercile in the regression. The decline in trips is smaller in magnitude for the richest areas by 0.265. Interestingly, driver pay increases by the most in the richest pickup areas, and fares don't increase proportionately more, leading to considerable increases in the labor share for the richest areas. The implied demand elasticity is -1.297/1.611 = -0.81 for the lowest income tercile, and (-1.297+.265)/(1.611-0.0674)=-0.67 for the higher income areas, i.e. demand is more inelastic in higher income areas.

Table 1c examines heterogeneity by access to public transit commuting patterns. Tercile 1 refers to census blocks where up to 24 percent commute by public transit, and Tercile 3 refers to census blocks where more than 51 percent commute by public transit. Demand is considerably more elastic in areas with high public transit usage ((-0.608-0.786)/(1.238+0.29)=-0.91), compared with low public transit usage (-0.608/1.238=-0.49), which may come from substitution to public transit.

## b. Driver-Level Regressions

**To be added in a future version of the draft.**

# 7. Preliminary Conclusions

We examined recent reforms in New York City that mandated per minute and per mile rates. We find that these reforms increased pay per trip. On average, these pay increases were passed on in the form of higher fares. We estimate that pay per trip and fares in NYC increased by about 15 percent due to the pay standard. By the law of demand, this will result in declines in trips. We find trips fall more in lower-income areas and areas with more access to public transit, the latter of which may indicate substitution to public transit. Overall, our estimates imply hourly wages increased by 4.5 percent on average. Increases in pay appear persistent over the 11 month period we study. We hope to provide more direct evidence on how the increases in pay per trip map to other important outcomes at the driver level, such as weekly take-home pay, in a future version of this draft.

## References

Abraham, Katharine, John Haltiwanger, Kristin Sandusky, James Spletzer. 2019. "Measuring the Gig Economy: Current Knowledge and Open Issues," in *Measuring and Accounting for Innovation in the 21st Century*. University of Chicago Press, 2019.

Angrist, Josh, Sydnee Caldwell and Jonathan Hall. 2018. "Uber vs. Taxi: A Driver's Eye View." NBER Working Paper #23891.

Chevalier, Judith, M. Keith Chen, Emily Oehlsen, and Peter E Rossi. "The Value of Flexible Work: Evidence from Uber Drivers." Forthcoming in the *Journal of Political Economy*.

Cohen, Peter, Robert Han, Jonathan Hall, Steven Levitt, and Robert Metcalfe. "Using Big Data to Estimate Consumer Surplus: The Case of Uber." NBER Working Paper #22627.

Collins, Brett, Andrew Garin, Emilie Jackson, and Mark Payne. 2019. "Is Gig Work Replacing Traditional Employment? Evidence from Two Decades of Tax Returns". *SOI Working Paper*.

Farber, Henry S. 2015. "Why you Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers." *Quarterly Journal of Economics*, 130, 4: 1975–2026.

Farrell, Diana, Fiona Greig, and Amar Hamoudi. 2019. "The Evolution of the Online Platform Economy: Evidence from Five Years of Banking Data." *AEA Papers and Proceedings*, 109: 362-66.

Hall, Jonathan and Alan Krueger. 2018. "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," *ILR Review*, 71(3): 705-732

Hall, Jonathan, John Horton and Daniel Knoepfle 2019. "Pricing Efficiently in Designed Markets: The Case of Ride-Sharing." *Working Paper*.

Koustas, Dmitri. 2019. "What Do Big Data Tell Us about Why People Take Gig Economy Jobs?" *AEA Papers and Proceedings*, 109: 367-71.

Koustas, Dmitri. 2018. "Consumption Insurance and Multijobholding: Evidence from Rideshare Drivers." *Mimeo*.

Mas, Alexandre and Amanda Pallais. 2018. "Valuing Alternative Work Arrangements." NBER WP.

Parrott, James and Michael Reich. 2018. "An Earnings Standard for New York City's App-based Drivers: Economic Analysis and Policy Assessment," Center for New York City Affairs at The New School, and Center on Wage and Employment Dynamics at University of California, Berkeley.

Figure 1: Total Trips, Average Driver Pay and Base Fares, and Total Drivers, August 2017-December 2019

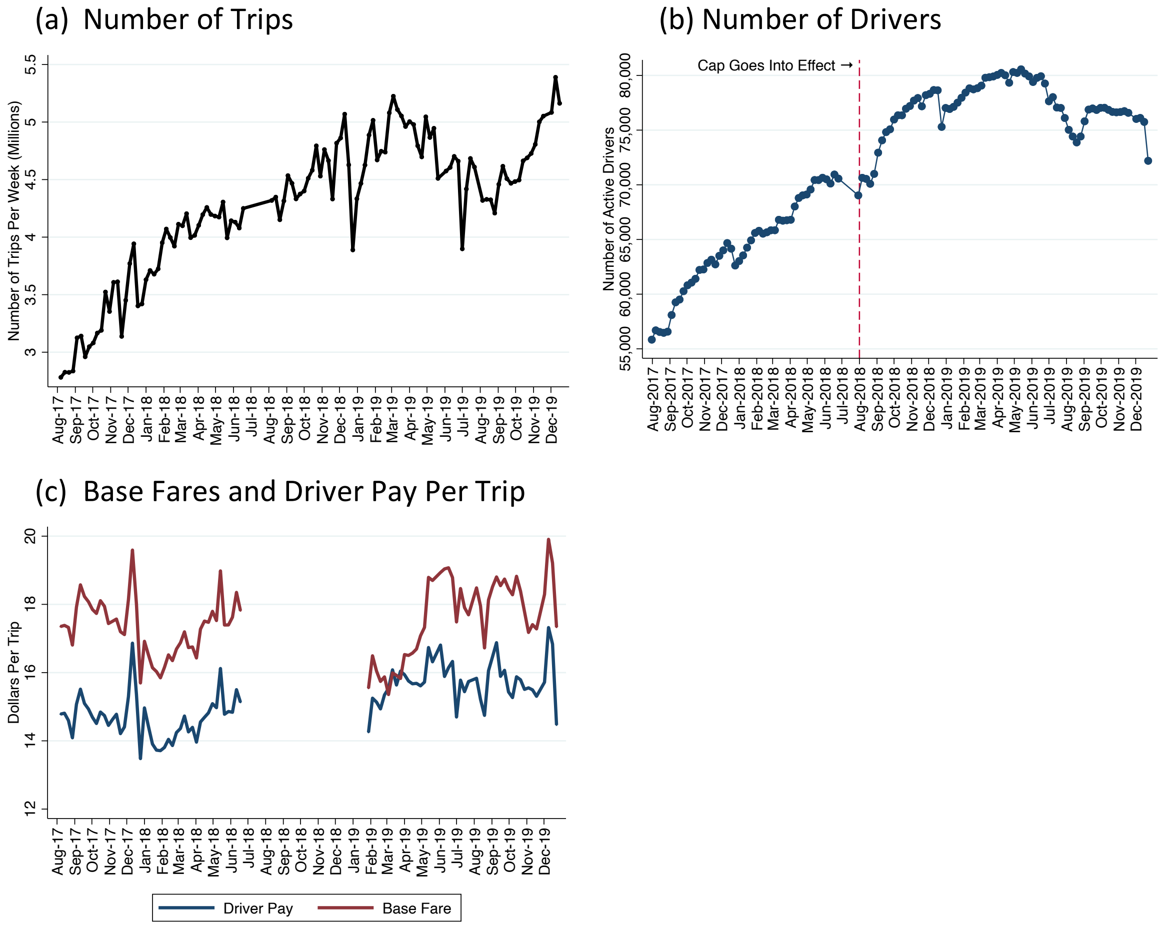


Notes: Panel (a) plots the total number of trips in millions. Panel (b) plots the total number of drivers active in the week. Panel (c) plots average base fares per trip and driver pay per trip. These are raw series, unadjusted for trip characteristics.

Figure 2: Variation in Pay, Trips and Hours Worked Across by Pickup Hour

(a) Average Driver Pay Per Trip

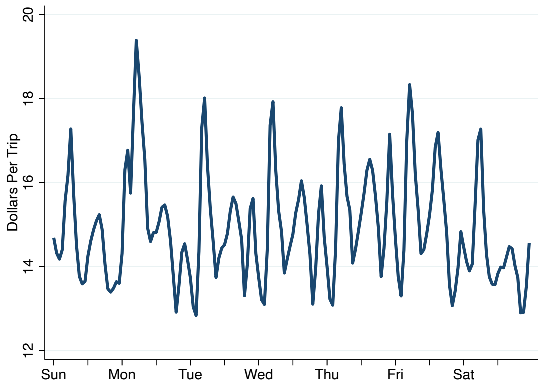


(b) Number of Trips and App Hours

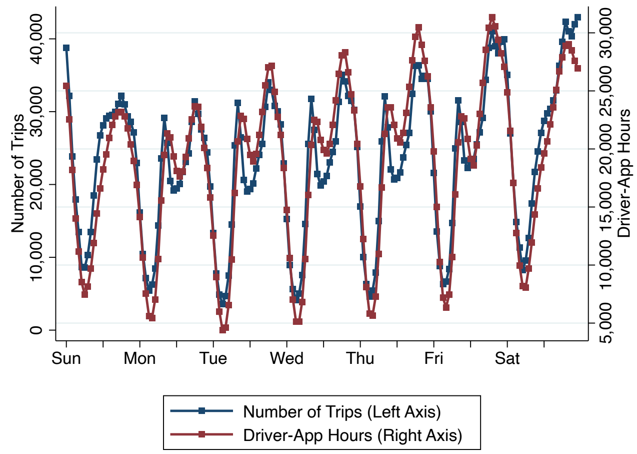


Notes: Panel (a) shows average driver pay per trip starting in the hour of day and day of week indicated (unadjusted for trip characteristics). Panel (b) shows the average number of trips for each hour of the day and day of the week (navy line) and the total number of hours drivers have at least one app on (maroon line). The averages are calculated over the period from July 2017-June 2018.

Figure 3: Differential Expected Exposure of the Pay Standard Across Hours and Days of the Week

(a) Share of Trips Counterfactual Pay Standard Expected to Bind

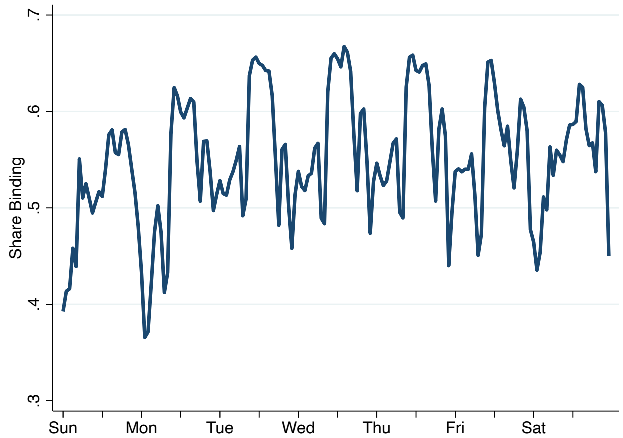


(b) Counterfactual Percent Increase in Pay Due to Pay Standard

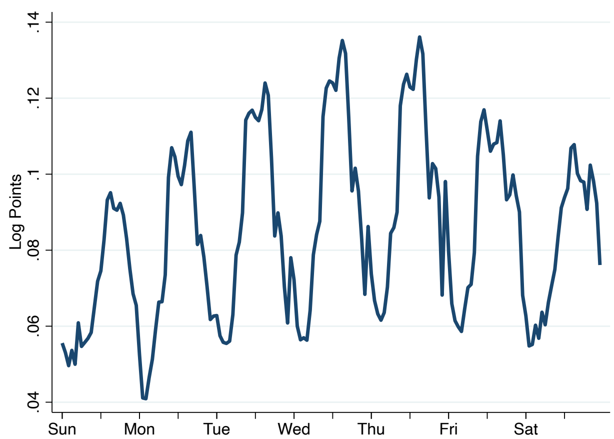




Notes: Panel (a) shows the share of trips for which a counterfactual pay standard would have binded. Panel (b) shows the counterfactual percentage increase in pay. Calculations are based on trips data from from July 2017-June 2018.

Figure 4: Counterfactual Percentage Increase in Driver Pay Due to Pay Standard, By Pickup Location

(a) by Pickup Location, Midnight

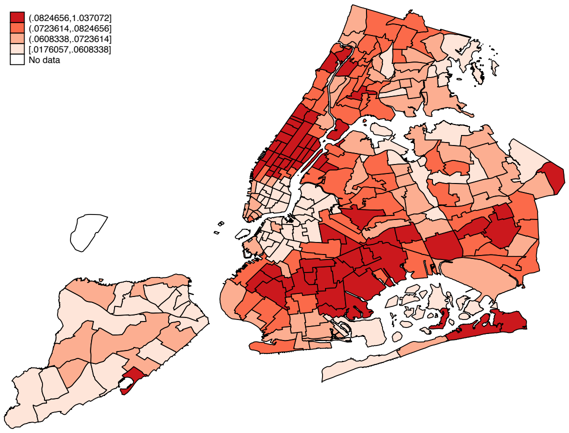


(b) by Dropoff Location, Midnight

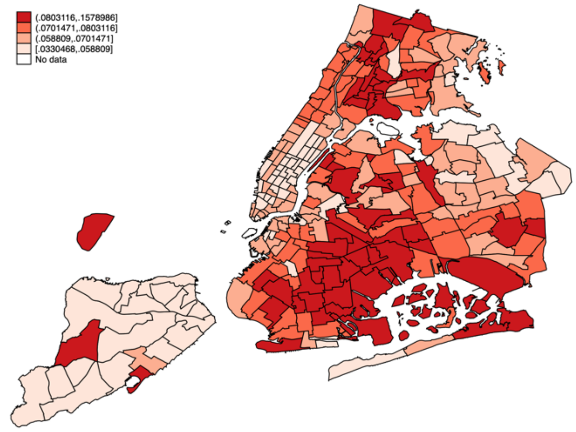


(c) by Pickup Location, Noon

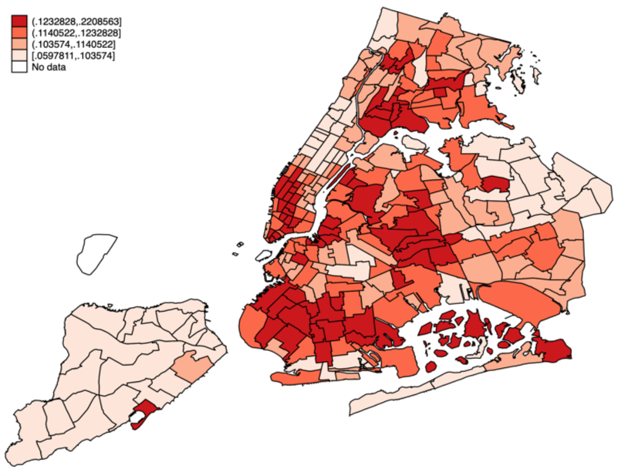


(d) by Dropoff Location, Noon

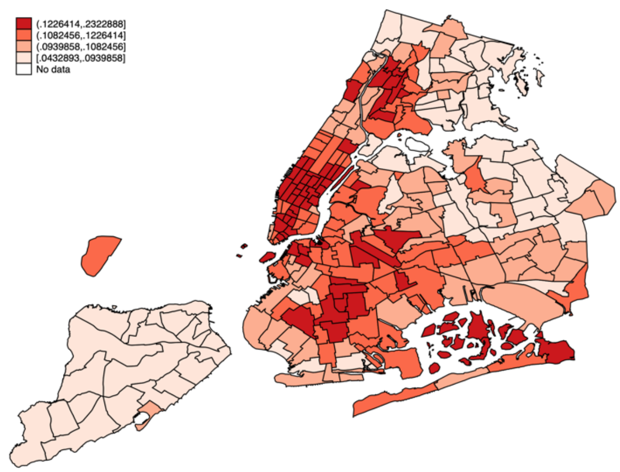


Notes: Figure shows variation in average expected increase in pay due to pay standard, by indicated hour, and the pickup or dropoff zone. Calculations are based on trips data from from July 2017-June 2018.

Figure 5: Main Results, Route-Level Regressions

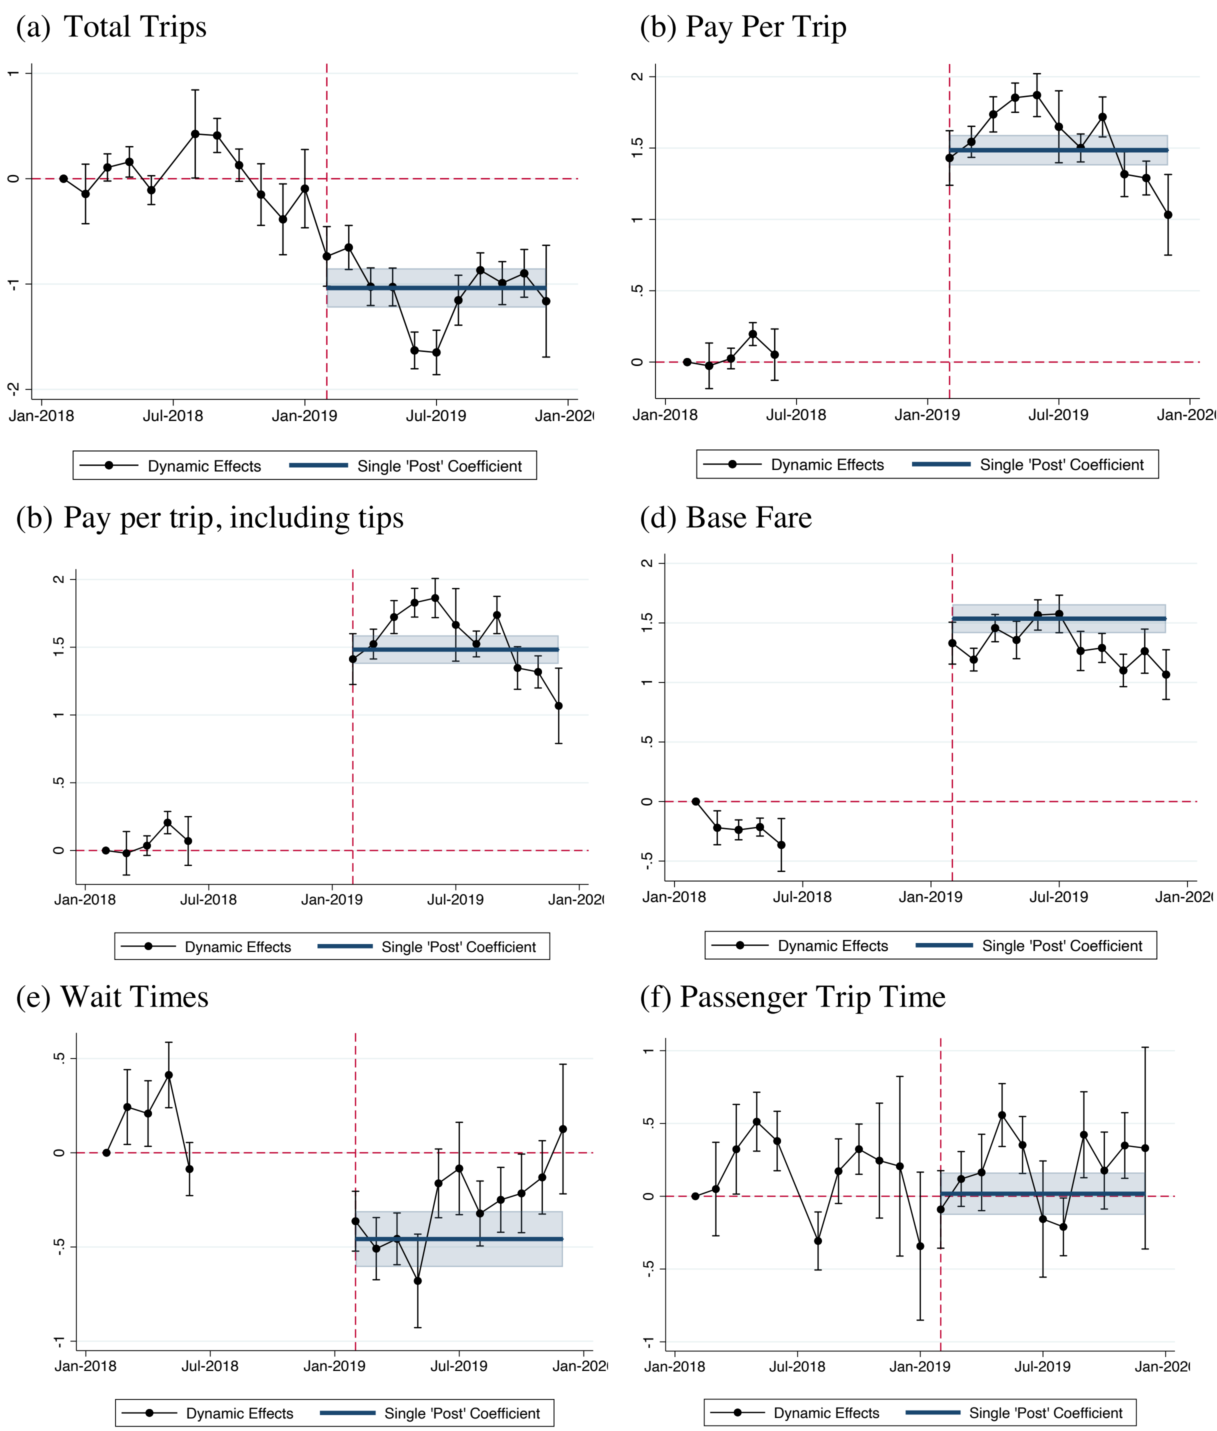


Notes: Figure plots estimated $\beta_j$ coefficients from Equation 4 (black dots) and $\beta$ coefficient from Equation 5 (blue line). $\beta_j$ coefficients are all relative to February 2018. 95% confidence bands are reported around the estimates. All outcomes are logged. The interpretation of the coefficients are elasticities of the indicated outcome with respect to the increase in pay from the pay standard.

Table 1a. Route-Level Regressions, Main Results, Single Post Coefficient

| | (1) N Trips | (2) Pay | (3) Pay+Tips | (4) Fare | (5) Labor Share | (6) Wait Time | (7) Trip Time | (8) Trip Miles |
|---|---|---|---|---|---|---|---|---|
| $Post \times \hat{pay}$ | -1.037*** | 1.485*** | 1.482*** | 1.536*** | -0.0421 | -0.458*** | 0.0176 | -0.0461 |
| | (0.0953) | (0.0547) | (0.0539) | (0.0621) | (0.0519) | (0.0757) | (0.0744) | (0.0352) |
| $Post \times CoreManhattan$ | -0.0718*** | -0.00969 | -0.00526 | -0.0396*** | 0.0208** | -0.01000 | 0.0636*** | 0.0720*** |
| | (0.0129) | (0.00870) | (0.00891) | (0.0102) | (0.00709) | (0.00947) | (0.00613) | (0.00266) |
| N Trips | 407,024,058 | 297,730,046 | 297,730,159 | 297,732,875 | 297,729,561 | 297,733,262 | 407,017,666 | 297,737,526 |
| N Routes | 19,918 | 19,902 | 19,902 | 19,902 | 19,902 | 19,902 | 19,918 | 19,903 |
| Dep. Pre' Mean | 3076.01 | 14.66 | 15.12 | 17.25 | .95 | 5.64 | 20.42 | 4.95 |

Notes: Table shows the regression results from estimating Equation 5. The regression includes week and route fixed effects. All outcomes are logged. The interpretation of the coefficients are elasticities of the indicated outcome with respect to the increase in pay from the pay standard. Standard errors two-way clustered on route and week in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 1b. Route-Level Regressions, Heterogeneity by Income Terciles

| | (1) N Trips | (2) Pay | (3) Pay+Tips | (4) Fare | (5) Labor Share | (6) Wait Time | (7) Trip Time | (8) Trip Miles |
|---|---|---|---|---|---|---|---|---|
| $Post \times \hat{pay}$ | -1.297*** | 1.186*** | 1.181*** | 1.611*** | -0.425*** | -0.626*** | 0.0714 | -0.117** |
| | (0.0759) | (0.0491) | (0.0485) | (0.0601) | (0.0531) | (0.0864) | (0.0903) | (0.0367) |
| $Post \times \hat{pay} \times tercile2$ | -0.0223 | 0.130*** | 0.133*** | -0.0389 | 0.159** | 0.145** | -0.0532* | 0.0648** |
| | (0.0554) | (0.0344) | (0.0342) | (0.0529) | (0.0467) | (0.0490) | (0.0252) | (0.0209) |
| $Post \times \hat{pay} \times tercile3$ | 0.265*** | 0.450*** | 0.454*** | -0.0674 | 0.535*** | 0.0183 | -0.105* | 0.0985*** |
| | (0.0519) | (0.0743) | (0.0733) | (0.0810) | (0.0640) | (0.0891) | (0.0416) | (0.0241) |
| $Post \times tercile2$ | 0.0118 | -0.0211*** | -0.0200*** | -0.00602 | -0.0142* | 0.0135* | 0.0195*** | 0.00425* |
| | (0.00654) | (0.00341) | (0.00341) | (0.00660) | (0.00605) | (0.00531) | (0.00290) | (0.00208) |
| $Post \times tercile3$ | -0.0622*** | -0.0553*** | -0.0515*** | -0.00470 | -0.0530*** | 0.00651 | 0.0441*** | 0.0105*** |
| | (0.00867) | (0.00751) | (0.00752) | (0.0101) | (0.00736) | (0.00901) | (0.00478) | (0.00288) |
| $Post \times CoreManhattan$ | -0.0675*** | -0.00485 | -0.00227 | -0.0384*** | 0.0251*** | -0.00545 | 0.0535*** | 0.0645*** |
| | (0.0104) | (0.00765) | (0.00784) | (0.00877) | (0.00576) | (0.00798) | (0.00555) | (0.00248) |
| N Trips | 370,696,111 | 284,356,645 | 284,357,044 | 284,363,821 | 284,355,122 | 284,352,203 | 370,660,508 | 284,377,454 |
| N Routes | 45,758 | 45,586 | 45,587 | 45,592 | 45,585 | 45,563 | 45,734 | 45,596 |
| Dep. Pre' Mean | 2096.45 | 14.04 | 14.47 | 16.4 | .95 | 5.59 | 19.74 | 4.64 |

Notes: Table shows the regression results from estimating Equation 6 where terciles refer to income terciles of the pickup Census tract. The regression includes week and route fixed effects. All outcomes are logged. The interpretation of the coefficients are elasticities of the indicated outcome with respect to the increase in pay from the pay standard. Standard errors two-way clustered on route and week in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 1c. Route-Level Regressions, Heterogeneity by Public-Transit-Share Terciles

| | (1) N Trips | (2) Pay | (3) Pay+Tips | (4) Fare | (5) Labor Share | (6) Wait Time | (7) Trip Time | (8) Trip Miles |
|---|---|---|---|---|---|---|---|---|
| $Post \times \hat{pay}$ | -0.608*** | 1.327*** | 1.305*** | 1.238*** | 0.0840 | -0.248** | 0.0497 | -0.148 |
| | (0.120) | (0.0616) | (0.0633) | (0.0945) | (0.0651) | (0.0791) | (0.0661) | (0.0978) |
| $Post \times \hat{pay} \times tercile2$ | -0.471*** | 0.264** | 0.297*** | 0.836*** | -0.551*** | -0.530*** | -0.163 | 0.0806 |
| | (0.132) | (0.0794) | (0.0816) | (0.107) | (0.0774) | (0.111) | (0.0881) | (0.0876) |
| $Post \times \hat{pay} \times tercile3$ | -0.786*** | 0.123 | 0.138 | 0.290** | -0.154 | -0.320*** | -0.0869 | 0.0534 |
| | (0.128) | (0.0711) | (0.0728) | (0.104) | (0.0772) | (0.0810) | (0.0764) | (0.0883) |
| $Post \times tercile2$ | 0.158*** | 0.000733 | -0.000851 | -0.124*** | 0.126*** | 0.0780*** | 0.0455*** | 0.00140 |
| | (0.0205) | (0.00880) | (0.00910) | (0.0128) | (0.00875) | (0.0135) | (0.0100) | (0.0113) |
| $Post \times tercile3$ | 0.184*** | 0.00706 | 0.00671 | -0.0563*** | 0.0634*** | 0.0447*** | 0.0199* | 0.00289 |
| | (0.0257) | (0.00796) | (0.00819) | (0.0113) | (0.00834) | (0.00989) | (0.00920) | (0.0116) |
| $Post \times CoreManhattan$ | -0.0620*** | -0.00928 | -0.00467 | -0.0396*** | 0.0212** | -0.00667 | 0.0700*** | 0.0750*** |
| | (0.0127) | (0.00883) | (0.00904) | (0.0103) | (0.00722) | (0.00949) | (0.00631) | (0.00276) |
| N Trips | 372,772,201 | 285,916,671 | 285,916,941 | 285,921,556 | 285,915,866 | 285,912,975 | 372,746,609 | 285,930,896 |
| N Routes | 31,369 | 31,111 | 31,111 | 31,117 | 31,110 | 31,082 | 31,321 | 31,131 |
| Dep. Pre' Mean | 2743.48 | 14.05 | 14.48 | 16.42 | .95 | 5.6 | 19.74 | 4.64 |

Notes: Table shows the regression results from estimating Equation 6 where terciles refer to terciles of the share of workers in the Census tract who commute to work via public transit. The regression includes week and route fixed effects. Standard errors two-way clustered on route and week in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.